

Complexity of the Two-Variable Fragment with Counting Quantifiers*

Ian Pratt-Hartmann

School of Computer Science, University of Manchester

March 21st, 2005

Abstract. The satisfiability and finite satisfiability problems for the two-variable fragment of first-order logic with counting quantifiers are both in NEXPTIME, even when counting quantifiers are coded succinctly.

Keywords: two-variable fragment, counting quantifiers, logic, complexity

1. Background

The two-variable fragment with counting quantifiers, here denoted \mathcal{C}^2 , is the set of function-free, first-order formulas containing at most two variables, but with the counting quantifiers $\exists_{\leq C}$, $\exists_{> C}$ and $\exists_{=C}$ (for every $C > 0$) allowed. The *satisfiability problem*, $\text{Sat-}\mathcal{C}^2$, is the problem of deciding, for a given formula ϕ of \mathcal{C}^2 , whether ϕ has a model; the *finite satisfiability problem*, $\text{Fin-Sat-}\mathcal{C}^2$, is the problem of deciding, for a given formula ϕ of \mathcal{C}^2 , whether ϕ has a finite model. It is well-known that \mathcal{C}^2 lacks the finite model property; hence $\text{Sat-}\mathcal{C}^2$ and $\text{Fin-Sat-}\mathcal{C}^2$ do not coincide. The decidability of $\text{Sat-}\mathcal{C}^2$ and $\text{Fin-Sat-}\mathcal{C}^2$ was shown by Grädel, Otto and Rosen [2]; the decidability of $\text{Sat-}\mathcal{C}^2$ was shown independently by Pacholski, Szwojda and Tendera [3, 4]. For a general survey, see Grädel and Otto [1].

When discussing the complexity of these problems, it is important to specify how the sizes of numerical quantifier subscripts are measured. Under *unary coding*, a quantifier subscript C is taken to have size C ; under *binary coding*, by contrast, the same subscript is taken to have size $\log C$. In determining upper complexity-bounds, binary coding is the more stringent accounting method, because formulas may be exponentially shorter under binary coding than they are under unary coding. Pacholski *et al.*, *op. cit.* showed that $\text{Sat-}\mathcal{C}^2$ is in NEXPTIME, but only under unary coding. The present paper shows that both $\text{Fin-Sat-}\mathcal{C}^2$ and $\text{Sat-}\mathcal{C}^2$ are in NEXPTIME, even under binary coding. It is well-known that the satisfiability problem for the two-variable fragment

* The author wishes to thank Dr Renate Schmidt and Dr Ulrike Sattler for their valuable help and suggestions.



without counting quantifiers (which has the finite model property) is NEXPTIME-hard. Hence, the bounds reported here are tight.

In the sequel, we confine attention to finite or countably infinite structures interpreting finite signatures of unary and binary predicates. Thus, all signatures are silently assumed to contain no individual constants or function-symbols. However, we treat the equality predicate \approx as a logical constant. The lack of individual constants constitutes no essential restriction of expressive power, because their effect can be reproduced using formulas of the form $\exists_{=1}xp(x)$. The presence of equality in the logic constitutes no extension of expressive power, because it can be defined by the formula $\forall x(x \approx x) \wedge \forall x\exists_{=1}y(x \approx y)$.

The structure of the paper is as follows. Section 2 establishes a normal form for \mathcal{C}^2 -formulas, and introduces various concepts which feature in the ensuing argument. Section 3 shows how any model of a normal-form \mathcal{C}^2 -formula ϕ can be recursively transformed into a model of ϕ in which, informally speaking, only a limited number of different sorts of element occur. Section 4 then outlines a method for encoding such models as data-structures satisfying certain arithmetical constraints, and shows on the basis of this encoding that $\text{Fin-Sat-}\mathcal{C}^2$ is in NEXPTIME. Finally, Section 5 extends this argument to show that $\text{Sat-}\mathcal{C}^2$ is also in NEXPTIME.

2. Preliminaries

If ϕ is a formula in \mathcal{C}^2 , let $\|\phi\|$ denote the number of symbols in ϕ , assuming binary coding of numerical quantifier subscripts.

Lemma 1. *Let ϕ be a formula in \mathcal{C}^2 . We can construct, in time bounded by a polynomial function of $\|\phi\|$, a formula*

$$\phi^* := \forall x\alpha \wedge \forall x\forall y(\beta \vee x \approx y) \wedge \bigwedge_{1 \leq h \leq m} \forall x\exists_{=C_h}y(f_h(x, y) \wedge x \not\approx y) \quad (1)$$

satisfying the conditions: (i) α is a quantifier-free, equality-free formula with x as its only variable; (ii) β is a quantifier-free, equality-free formula with x and y as its only variables; (iii) m is a positive integer; (iv) for all h ($1 \leq h \leq m$), f_h is a binary predicate and C_h a positive integer; and (v) for any domain A of size greater than $C = \max_{1 \leq h \leq m} C_h$, ϕ^ is satisfiable over A if and only if ϕ is satisfiable over A .*

Proof. Routine adaptation of textbook transformation to Scott normal form. See, for example, Grädel and Otto [1] Sec. 2.1 for an explanation of the required techniques. \square

With binary coding of quantifier subscripts, the quantity $C = \max_{1 \leq h \leq m} C_h$ in Lemma 1 satisfies $C \leq 2^{\|\phi^*\|}$. Hence, the problem of determining whether ϕ is satisfiable over a domain of size C or less is in NEXPTIME. It thus follows from Lemma 1 that, to show that Fin-Sat- \mathcal{C}^2 and Sat- \mathcal{C}^2 are in NEXPTIME, we may restrict attention to formulas of the form (1).

Fix ϕ^* to be some formula of the form (1), and let Σ^* be the signature of ϕ^* . Thus, $\alpha(x)$ is a quantifier-free, equality-free formula over Σ^* with x as its only variable, $\beta(x, y)$ is a quantifier-free, equality-free formula over Σ^* with x and y as its only variables, f_1, \dots, f_m are distinct binary predicates of Σ^* (with $m > 0$), and C_1, \dots, C_m are positive integers. In addition, let $C = \max_{1 \leq h \leq m} C_h$. We keep the meanings of these symbols unchanged throughout this paper. The predicates f_1, \dots, f_m will play a key role in the ensuing argument; we refer to these predicates as the *counting predicates*. In addition, let us say that a signature Σ is a *unary extension* of Σ^* if Σ consists of Σ^* together with a (possibly empty) finite set of new unary predicates.

We review some standard concepts. A *literal* is an atomic formula or the negation of an atomic formula. Let Σ be a finite signature (consisting of unary and binary predicates only). A *1-type* (over Σ) is a maximal consistent set of equality-free literals over Σ involving only the variable x . A *2-type* (over Σ) is a maximal consistent set of equality-free literals over Σ involving only the variables x and y . Reference to Σ is suppressed where clear from context. If τ is a 2-type, then the result of transposing the variables x and y in τ will also be a 2-type, denoted τ^{-1} . If \mathfrak{A} is any structure interpreting Σ , and $a \in A$, then there exists a unique 1-type $\pi(x)$ over Σ such that $\mathfrak{A} \models \pi[a]$; we denote π by $\text{tp}^{\mathfrak{A}}[a]$. If, in addition, $b \in A$ is distinct from a , then there exists a unique 2-type $\tau(x, y)$ over Σ such that $\mathfrak{A} \models \tau[a, b]$; we denote τ by $\text{tp}^{\mathfrak{A}}[a, b]$. We do not define $\text{tp}^{\mathfrak{A}}[a, b]$ if $a = b$. If π is a 1-type, we say that π is *realized* in \mathfrak{A} if there exists $a \in A$ with $\text{tp}^{\mathfrak{A}}[a] = \pi$. If τ is a 2-type, we say that τ is *realized* in \mathfrak{A} if there exist distinct $a, b \in A$ with $\text{tp}^{\mathfrak{A}}[a, b] = \tau$.

Notation 1. *Given a fixed signature Σ , any 2-type τ includes a unique 1-type, denoted $\text{tp}_1(\tau)$; in addition, we write $\text{tp}_2(\tau)$ for $\text{tp}_1(\tau^{-1})$.*

Remark 1. *If $\text{tp}^{\mathfrak{A}}[a, b] = \tau$, then $\text{tp}^{\mathfrak{A}}[b, a] = \tau^{-1}$, $\text{tp}^{\mathfrak{A}}[a] = \text{tp}_1(\tau)$ and $\text{tp}^{\mathfrak{A}}[b] = \text{tp}_2(\tau)$.*

Definition 1. Let Σ be a unary extension of Σ^* , and let τ be a 2-type over Σ . We say that τ is a *message-type* (over Σ) if, for some counting predicate f_h ($1 \leq h \leq m$), $f_h(x, y) \in \tau$. If τ is a message-type such that τ^{-1} is also a message-type, we say that τ is *invertible*. On the other hand, if τ is a 2-type such that neither τ nor τ^{-1} is a message-type, we say that τ is *silent*.

Thus, a 2-type τ is an invertible message-type if and only if there are counting predicates f_h and $f_{h'}$ ($1 \leq h \leq m$, $1 \leq h' \leq m$) such that $f_h(x, y) \in \tau$ and $f_{h'}(y, x) \in \tau$. The terminology is meant to suggest the following imagery. If $\text{tp}^{\mathfrak{A}}[a, b]$ is a message-type μ , then we may imagine that a sends a message (of type μ) to b . If μ is invertible, then b replies by sending a message (of type μ^{-1}) back to a . If $\text{tp}^{\mathfrak{A}}[a, b]$ is silent, then neither element sends a message to the other.

The remainder of this section is devoted to some auxiliary observations regarding models of ϕ^* interpreting unary extensions of Σ^* . In particular, we introduce the concepts of *chromaticity* and *differentiation* for such structures.

Definition 2. Let Σ be a unary extension of Σ^* , and let \mathfrak{A} be a structure interpreting Σ . We say that \mathfrak{A} is *chromatic* (over Σ) if distinct elements connected by a chain of 1 or 2 invertible message-types have distinct 1-types. That is, \mathfrak{A} is chromatic just in case, for all $a, a', a'' \in A$:

1. if $a \neq a'$ and $\text{tp}^{\mathfrak{A}}[a, a']$ is an invertible message-type, then $\text{tp}^{\mathfrak{A}}[a] \neq \text{tp}^{\mathfrak{A}}[a']$; and
2. if a, a', a'' are pairwise distinct and both $\text{tp}^{\mathfrak{A}}[a, a']$ and $\text{tp}^{\mathfrak{A}}[a', a'']$ are invertible message-types, then $\text{tp}^{\mathfrak{A}}[a] \neq \text{tp}^{\mathfrak{A}}[a'']$.

Any model of ϕ^* can be made chromatic by interpreting not-too-many new unary predicates:

Lemma 2. *Let the structure \mathfrak{A} interpret Σ^* , and let Σ' be the signature formed by adding $\log((mC)^2 + 1)$ (rounded up) new unary predicates to Σ^* . If $\mathfrak{A} \models \phi^*$, then \mathfrak{A} can be expanded to a chromatic structure \mathfrak{A}' interpreting Σ' .*

Proof. Suppose $\mathfrak{A} \models \phi^*$, and consider the (undirected) graph G on A whose edges are the pairs of distinct elements connected by a chain of 1 or 2 invertible message-types. That is, $G = (A, E^1 \cup E^2)$, where

$$\begin{aligned} E^1 &= \{(a, a') \mid a \neq a' \text{ and } \text{tp}^{\mathfrak{A}}[a, a'] \text{ is an invertible message-type}\} \\ E^2 &= \{(a, a'') \mid a \neq a'' \text{ and for some } a' \in A, (a, a') \text{ and } (a', a'') \text{ are} \\ &\quad \text{both in } E^1\}. \end{aligned}$$

Since $\mathfrak{A} \models \phi^*$, the degree of G (in the normal graph-theoretic sense) is at most $(mC)^2$. Now use the standard (greedy) algorithm to colour the nodes of G with $(mC)^2 + 1$ colours in such a way that no edge joins two nodes of the same colour. By interpreting the $\log((mC)^2 + 1)$ (rounded up) new unary predicates to encode these colours, we obtain the desired expansion \mathfrak{A}' . \square

Chromatic models will play an important part in the argument of Section 4. They are easy to work with because they exhibit the following properties:

Remark 2. *Let Σ be a unary extension of Σ^* and let π be a 1-type over Σ . Suppose \mathfrak{A} is a chromatic structure interpreting Σ , and let $a \in A$. Then there is at most one element $a' \in A \setminus \{a\}$ with 1-type π such that a sends an invertible message to a' ; moreover, if the 1-type of a is itself π , there is no such element a' .*

Turning now to the concept of differentiation, fix the constant Z to be $(mC + 1)^2$. The reasons for our particular choice of Z will become clear in the course of the paper.

Definition 3. Let \mathfrak{A} be a structure interpreting a signature Σ . We say that \mathfrak{A} is *differentiated (over Σ)* if, for every 1-type π over Σ , the number u of elements in A having 1-type π satisfies either $u \leq 1$ or $u > Z$.

Thus, in a differentiated structure, every 1-type is realized either at most once or more than Z times. Any structure can be made differentiated by interpreting not-too-many new unary predicates:

Lemma 3. *Let Σ be a unary extension of Σ^* , and let \mathfrak{A} be a structure interpreting Σ . Let Σ' be the signature formed by adding $\log Z$ (rounded up) new unary predicates to Σ (so that Σ' is also a unary extension of Σ^*). Then \mathfrak{A} can be expanded to a differentiated structure \mathfrak{A}' interpreting Σ' . Moreover, if \mathfrak{A} is chromatic over Σ , then \mathfrak{A}' is chromatic over Σ' .*

Proof. For each 1-type π realized more than once but no more than Z times, colour the elements having 1-type π using Z different colours. By interpreting the $\log Z$ (rounded up) new unary predicates to encode these colours, we obtain the desired expansion \mathfrak{A}' . This process clearly preserves chromaticity. \square

Now let us fix a signature Σ obtained by adding $2 \log Z$ (rounded up) new unary predicates to Σ^* . Since $Z > (mC)^2 + 1$, Lemmas 2 and 3 guarantee that, if ϕ^* has a model at all, then it has a model with the same domain which is chromatic and differentiated over Σ . Denote the total number of symbols in Σ by s . Enumerate the 1-types over Σ , in some arbitrary order, as π_1, \dots, π_L ; enumerate the invertible message-types over Σ , in some arbitrary order, as μ_1, \dots, μ_{M^*} ; and enumerate the non-invertible message-types over Σ , in some arbitrary order, as $\mu_{M^*+1}, \dots, \mu_M$. (Thus, μ_1, \dots, μ_M is an enumeration of all

Table I. Quick reference guide to symbols with fixed interpretations.

ϕ^*	$\forall x\alpha \wedge \forall x\forall y(\beta \vee x \approx y) \wedge \bigwedge_{1 \leq h \leq m} \forall x \exists_{=C_h} y (f_h(x, y) \wedge x \not\approx y)$
C	$\max_{1 \leq h \leq m} C_h$
Z	$(mC + 1)^2$
Σ^*	signature of ϕ^*
Σ	Σ^* with $2 \log Z$ (rounded up) new unary predicates
s	the number of symbols in Σ
π_1, \dots, π_L	the 1-types over Σ
μ_1, \dots, μ_{M^*}	the invertible message-types over Σ
$\mu_{M^*+1}, \dots, \mu_M$	the non-invertible message-types over Σ
Ξ	the set of silent 2-types over Σ

the message-types over Σ .) Finally, denote the set of silent 2-types over Σ by Ξ . We fix the symbols Σ , s , L , π_i ($1 \leq i \leq L$), M^* , M , μ_j ($1 \leq j \leq M$) and Ξ to have these meanings for the remainder of the paper. Since Σ is the only signature we shall be concerned with in the sequel, we generally suppress reference to it. Thus, ‘model of ϕ^* ’ henceforth means ‘model of ϕ^* interpreting Σ ’, ‘1-type’ means ‘1-type over Σ ’, and so on. Table I lists the symbols introduced in this section and their fixed interpretations. Evidently, m and s are bounded by a polynomial function of $\|\phi^*\|$; moreover, C , Z , L , M^* and M are bounded by an exponential function of $\|\phi^*\|$. Thus, we should regard m and s as ‘small’, and C , Z , L , M^* and M as ‘large’.

Using the notational conventions just established, we can state a simple fact about differentiated models of ϕ^* which will prove useful at several points in the sequel.

Definition 4. Let \mathfrak{A} be a structure and π, π' 1-types. We say that π and π' form a *noisy pair* in \mathfrak{A} if, for all distinct $a, a' \in A$ such that $\text{tp}^{\mathfrak{A}}[a] = \pi$ and $\text{tp}^{\mathfrak{A}}[a'] = \pi'$, either $\text{tp}^{\mathfrak{A}}[a, a']$ or $\text{tp}^{\mathfrak{A}}[a', a]$ is a message-type.

Informally, π and π' form a noisy pair just in case every element with 1-type π either sends a message to, or receives a message from, every element (itself excepted) with 1-type π' . Note that Definition 4 does not require π and π' to be distinct.

Lemma 4. *Suppose that \mathfrak{A} is a differentiated model of ϕ^* , and that the 1-types π and π' form a noisy pair in \mathfrak{A} . Then either there is at most 1 element of A having 1-type π , or there is at most 1 element of A having 1-type π' .*

Proof. Suppose for contradiction that π and π' form a noisy pair, and that there is more than one element having 1-type π and more than one

element having 1-type π' . Since \mathfrak{A} is differentiated, there are at least $(mC + 1)^2 + 1$ elements having 1-type π and at least $(mC + 1)^2 + 1$ elements having 1-type π' . Now let B be a set of elements having 1-type π , with $|B| = (mC)^2 + mC + 1$; and let B' be a set of elements having 1-type π' , disjoint from B , with $|B'| = mC + 1$. Since $|B| + |B'| = (mC + 1)^2 + 1$, such sets can evidently be found, even if $\pi = \pi'$. Select any $b \in B$. Since $\mathfrak{A} \models \phi^*$, b sends a message to at most mC elements in B' , so that there exists $b' \in B'$ such that b sends no message to b' and hence (since π and π' form a noisy pair) such that b' sends a message to b . Thus, for all $b \in B$, there exists $b' \in B'$ such that b' sends a message to b . But since each of the $(mC + 1)$ elements of $|B'|$ sends a message to at most mC elements in B , we have $|B| \leq mC(mC + 1)$, contradicting $|B| = (mC)^2 + mC + 1$. \square

3. Manipulating structures

The goal of this section is to show that, if ϕ^* has a model, then it has a model in which only a limited number of different sorts of element occur, in a sense which we must make precise. We remind ourselves that the symbols ϕ^* , α , β , m , C_h , f_h ($1 \leq h \leq m$), C , Z , Σ , s , M^* , M , μ_j ($1 \leq j \leq M$) have fixed interpretations (Table I), and that the signature Σ is assumed throughout.

3.1. PROFILES, COUNTS AND APPROXIMATIONS

The first step is to manufacture some tools for manipulating structures.

Notation 2. Let \mathfrak{A} be a structure, π a 1-type, and Π a set of 1-types. When \mathfrak{A} is clear from context, denote by A_π the set $\{a \in A \mid \text{tp}^{\mathfrak{A}}[a] = \pi\}$, and denote by A_Π the set $\{a \in A \mid \text{tp}^{\mathfrak{A}}[a] \in \Pi\}$. In addition, denote by Π^c the set of all and only those 1-types not contained in Π .

Remark 3. For any structure \mathfrak{A} and any set of 1-types Π , $A_{\Pi^c} = A \setminus A_\Pi$.

For the next definition, recall that μ_1, \dots, μ_M are the message-types (invertible and non-invertible).

Definition 5. Suppose $\mathfrak{A} \models \phi^*$. Let $a \in A$, and let Π be any set of 1-types. The Π -profile of a in \mathfrak{A} , denoted $\text{pr}_\Pi^{\mathfrak{A}}[a]$, is the M -element integer vector whose j th element ($1 \leq j \leq M$) is given by:

$$|\{b \in A_\Pi : b \neq a \text{ and } \text{tp}^{\mathfrak{A}}[a, b] = \mu_j\}|.$$

If Π is the set of all 1-types, we call $\text{pr}_\Pi^{\mathfrak{A}}[a]$ simply the *profile* of a in \mathfrak{A} , and denote it $\text{pr}^{\mathfrak{A}}[a]$.

The vector $\text{pr}^{\mathfrak{A}}[a]$ records, for each message-type μ_j ($1 \leq j \leq M$), how many elements a sends a message of type μ_j to. For any set of 1-types Π , the vector $\text{pr}_{\Pi}^{\mathfrak{A}}[a]$ is an incomplete description of $\text{pr}^{\mathfrak{A}}[a]$ obtained by zeroing its j th coordinate whenever $\text{tp}_2(\mu_j)$ is not a member of Π . It helps to think of these vectors as describing aspects of a 's 'local environment'.

For the next definition, recall that f_1, \dots, f_m are the counting predicates.

Definition 6. Suppose $\mathfrak{A} \models \phi^*$. Let $a \in A$, and let Π be any set of 1-types. The Π -count of a in \mathfrak{A} , denoted $\text{ct}_{\Pi}^{\mathfrak{A}}[a]$, is the m -element integer vector whose h th element ($1 \leq h \leq m$) is given by:

$$|\{b \in A_{\Pi} : b \neq a \text{ and } \mathfrak{A} \models f_h[a, b]\}|.$$

If Π is the set of all 1-types, we call $\text{ct}_{\Pi}^{\mathfrak{A}}[a]$ simply the *count* of a in \mathfrak{A} , and denote it $\text{ct}^{\mathfrak{A}}[a]$.

The vector $\text{ct}^{\mathfrak{A}}[a]$ records, for each counting predicate f_h ($1 \leq h \leq m$), how many elements a is non-reflexively related to by f_h . For any set of 1-types Π , the vector $\text{ct}_{\Pi}^{\mathfrak{A}}[a]$ is an incomplete description of $\text{ct}^{\mathfrak{A}}[a]$ obtained by discounting those elements whose 1-type is not in Π . It is best to think of the vector $\text{ct}_{\Pi}^{\mathfrak{A}}[a]$ as providing a statistical summary of the vector $\text{pr}_{\Pi}^{\mathfrak{A}}[a]$. In particular, $\text{pr}_{\Pi}^{\mathfrak{A}}[a]$ determines $\text{ct}_{\Pi}^{\mathfrak{A}}[a]$, but not conversely.

Suppose $\mathfrak{A} \models \phi^*$, and let $a \in A$. Then $\text{pr}^{\mathfrak{A}}[a]$ is a vector of M integers in the range $[0, C]$; hence, the number of different profile vectors realized in \mathfrak{A} is bounded by $(C + 1)^M$ and therefore by a doubly exponential function of $\|\phi^*\|$. On the other hand, $\text{ct}^{\mathfrak{A}}[a]$ is a vector of m integers in the range $[0, C]$; hence, the number of different count vectors realized in \mathfrak{A} is bounded by $(C + 1)^m$, and therefore by a singly exponential function of $\|\phi^*\|$. The main task of this section is to show that, by modifying the structure \mathfrak{A} , the number of realized profile vectors can be reduced so that it too is bounded by a singly exponential function of $\|\phi^*\|$. The next definition introduces the device used to effect this reduction.

Definition 7. Suppose \mathfrak{A} is a chromatic model of ϕ^* . Let Π be a set of 1-types, and let B be a subset of A . A structure \mathfrak{A}' over the domain A is a (Π, B) -approximation to \mathfrak{A} if (i) \mathfrak{A}' is chromatic; (ii) every 2-type realized in \mathfrak{A}' is also realized in \mathfrak{A} ; and (iii) for all $a \in A$:

1. $\text{tp}^{\mathfrak{A}'}[a] = \text{tp}^{\mathfrak{A}}[a]$;
2. $\text{pr}_{\Pi^c}^{\mathfrak{A}'}[a] = \text{pr}_{\Pi^c}^{\mathfrak{A}}[a]$;

3. $a \in A \setminus B$ implies $\text{pr}^{\mathfrak{A}'}[a] = \text{pr}^{\mathfrak{A}}[a]$;
4. $a \in B$ implies $\text{ct}_{\Pi}^{\mathfrak{A}'}[a] = \text{ct}_{\Pi}^{\mathfrak{A}}[a]$.

Very roughly, a (Π, B) -approximation to \mathfrak{A} is a surgically modified version of \mathfrak{A} in which only the Π -profiles of elements of B have been interfered with. In particular: all elements of A retain their old 1-types and their old Π^c -profiles; all elements of $A \setminus B$ retain their old profiles; and all elements of B retain their old Π -counts. In addition, chromaticity is preserved, and no new 2-types (or 1-types) are introduced. We remark that, in Condition 4 of Definition 7 (iii), the restriction that $a \in B$ is in fact logically redundant, since if $a \notin B$, Condition 3 certainly entails $\text{ct}_{\Pi}^{\mathfrak{A}'}[a] = \text{ct}_{\Pi}^{\mathfrak{A}}[a]$.

Remark 4. *Let \mathfrak{A} , \mathfrak{A}' and \mathfrak{A}'' be chromatic models of ϕ^* over some common domain A , let Π , Π' be sets of 1-types and let B , B' be subsets of A . Then \mathfrak{A} is a (Π, B) -approximation to itself. Furthermore, if \mathfrak{A}' is a (Π, B) -approximation to \mathfrak{A} , $\Pi \subseteq \Pi'$, and $B \subseteq B'$, then \mathfrak{A}' is also a (Π', B') -approximation to \mathfrak{A} . Finally, if \mathfrak{A}' is a (Π, B) -approximation to \mathfrak{A} , and \mathfrak{A}'' is a (Π, B) -approximation to \mathfrak{A}' , then \mathfrak{A}'' is a (Π, B) -approximation to \mathfrak{A} .*

A crucial fact about (Π, B) -approximations is that they maintain satisfaction of ϕ^* :

Lemma 5. *Suppose \mathfrak{A} is a chromatic model of ϕ^* . Let Π be a set of 1-types, let B be a subset of A , and let \mathfrak{A}' be a (Π, B) -approximation to \mathfrak{A} . Then $\mathfrak{A}' \models \phi^*$.*

Proof. By Remark 4, we may assume without loss of generality that Π is the set of all 1-types and $B = A$. Since every 2-type realized in \mathfrak{A}' is also realized in \mathfrak{A} , $\mathfrak{A}' \models \forall x \alpha \wedge \forall x \forall y (\beta \vee x \approx y)$. And since $\text{ct}^{\mathfrak{A}'}[a] = \text{ct}^{\mathfrak{A}}[a]$ for every $a \in A$, $\mathfrak{A}' \models \bigwedge_{1 \leq h \leq m} \forall x \exists =_{C_h} y (f_h(x, y) \wedge x \not\approx y)$. \square

3.2. GROUPS AND PATCHES

Lemma 5 shows that, for any set of 1-types Π and any $B \subseteq A$, taking a (Π, B) -approximation to a model \mathfrak{A} of ϕ^* yields another model of ϕ^* . Our strategy now is to show that, for certain Π and B , a (Π, B) -approximation can be obtained in which the elements of B exhibit ‘few’ Π -profiles.

For the next definition, recall that μ_1, \dots, μ_{M^*} are the *invertible* message-types. Thus, for any model \mathfrak{A} of ϕ^* and any 1-type π , the first M^* coordinates of any $\{\pi\}$ -profile $\text{pr}_{\{\pi\}}^{\mathfrak{A}}[a]$ tell us, for each invertible

message-type μ_j ($1 \leq j \leq M^*$), how many elements in A_π the element a sends a message of type μ_j to. Notice that, if \mathfrak{A} is chromatic, then, by Remark 2, the first M^* coordinates of $\text{pr}_{\{\pi\}}^{\mathfrak{A}}[a]$ are either all zero, or else are all zero except for a single occurrence of unity.

Definition 8. Suppose $\mathfrak{A} \models \phi^*$. Let Π be a set of 1-types, let π be a 1-type, and let B be a subset of A . We say that B is a Π -group if every element of B has the same 1-type and every element of B has the same Π -count. We say that B is a π -patch if B is a $\{\pi\}$ -group and, for all $a, b \in B$, the vectors $\text{pr}_{\{\pi\}}^{\mathfrak{A}}[a]$ and $\text{pr}_{\{\pi\}}^{\mathfrak{A}}[b]$ agree in each of their first M^* coordinates.

We now demonstrate that, if B is a π -patch in a chromatic, differentiated model of ϕ^* , then by taking a $(\{\pi\}, B)$ -approximation to that model, we can reduce the number of $\{\pi\}$ -profiles realized by the elements of B to just 1.

Lemma 6. *Suppose \mathfrak{A} is a chromatic, differentiated model of ϕ^* . Let π be a 1-type, and let $B \subseteq A$ be a π -patch in \mathfrak{A} . Then there exists a structure \mathfrak{A}' such that \mathfrak{A}' is a $(\{\pi\}, B)$ -approximation to \mathfrak{A} in which the elements of B all have the same $\{\pi\}$ -profile.*

Proof. Let π^* be the 1-type such that $B \subseteq A_{\pi^*}$. If $|B| \leq 1$, $\mathfrak{A}' = \mathfrak{A}$ obviously satisfies the conditions of the Lemma. And if $|A_\pi| \leq 1$, by re-interpreting the non-counting predicates if necessary, we easily obtain a structure satisfying the conditions of the lemma. So we may suppose that B and A_π both contain more than one element. Since \mathfrak{A} is differentiated, by Lemma 4, let τ be a silent 2-type such that, for some $a \in A_{\pi^*}$ and some $a' \in A_\pi$, $\text{tp}^{\mathfrak{A}}[a, a'] = \tau$.

For $a \in B$, let

$$A_a = \{a' \in A_\pi \mid a \neq a' \text{ and } \text{tp}^{\mathfrak{A}}[a, a'] \text{ is a non-invertible message-type}\};$$

and for $a \notin B$, let $A_a = \emptyset$. Notice, incidentally, that $a' \in A_a$ implies $a \notin A_{a'}$. Choose $b \in B$ for which $|A_b|$ is smallest, and fix b . Enumerate A_b as b_1, b_2, \dots . For any $a \in B$ not equal to b , let \hat{A}_a be a subset of A_a having the same number of elements as A_b , and enumerate \hat{A}_a as a_1, a_2, \dots .

We now define the structure \mathfrak{A}' by assigning 2-types as follows. For all $a \in B$ such that $a \neq b$, set

$$\text{tp}^{\mathfrak{A}'}[a, a_i] = \text{tp}^{\mathfrak{A}}[b, b_i], \quad (2)$$

where i ranges over the enumeration of \hat{A}_a , and set

$$\text{tp}^{\mathfrak{A}'}[a, a'] = \tau, \quad (3)$$

where a' is any element of $A_a \setminus \hat{A}_a$. In addition, for all distinct a, a' such that $a' \notin A_a$ and $a \notin A_{a'}$, set

$$\text{tp}^{\mathfrak{A}'}[a, a'] = \text{tp}^{\mathfrak{A}}[a, a']. \quad (4)$$

Since $a' \in A_a$ implies $a \notin A_{a'}$, none of these assignments overwrites any other. And since $B \subseteq A_{\pi^*}$, the 1-type assignments implicit in (2)–(4) never clash: indeed, we have $\text{tp}^{\mathfrak{A}'}[a] = \text{tp}^{\mathfrak{A}}[a]$ for all $a \in A$. Furthermore, the transformation from \mathfrak{A} to \mathfrak{A}' does not affect invertible message-types. That is: for distinct a, a' , $\text{tp}^{\mathfrak{A}}[a, a']$ is an invertible message-type if and only if $\text{tp}^{\mathfrak{A}'}[a, a']$ is an invertible message-type; and moreover, if $\text{tp}^{\mathfrak{A}}[a, a']$ is an invertible message-type, then $\text{tp}^{\mathfrak{A}'}[a, a'] = \text{tp}^{\mathfrak{A}}[a, a']$.

We now verify that \mathfrak{A}' is a $(\{\pi\}, B)$ -approximation to \mathfrak{A} . From the remarks of the previous paragraph and the fact that \mathfrak{A} is chromatic, we have that \mathfrak{A}' is also chromatic. In addition, it is immediate from (2)–(4) that every 2-type realized in \mathfrak{A}' is also realized in \mathfrak{A} . Now let a be any element of B . Since B is a π -patch, the vectors $\text{pr}_{\{\pi\}}^{\mathfrak{A}}[a]$ and $\text{pr}_{\{\pi\}}^{\mathfrak{A}'}[a]$ by definition agree in their first M^* coordinates (corresponding to the invertible message-types). Hence, since we have just shown that the transformation from \mathfrak{A} to \mathfrak{A}' does not affect invertible message-types, the vectors $\text{pr}_{\{\pi\}}^{\mathfrak{A}'}[a]$ and $\text{pr}_{\{\pi\}}^{\mathfrak{A}}[a]$ also agree in their first M^* coordinates. Furthermore, the assignments (2)–(4) guarantee that $\text{pr}_{\{\pi\}}^{\mathfrak{A}'}[a]$ and $\text{pr}_{\{\pi\}}^{\mathfrak{A}}[a]$ also agree in the remaining coordinates $M^* + 1, \dots, M$ (corresponding to the non-invertible message-types). Hence,

$$\text{pr}_{\{\pi\}}^{\mathfrak{A}'}[a] = \text{pr}_{\{\pi\}}^{\mathfrak{A}}[a] \quad \text{for all } a \in B. \quad (5)$$

It is now a simple matter to check the numbered conditions in Definition 7 (iii). Let a be an arbitrary element of A .

1. We have already established that $\text{tp}^{\mathfrak{A}'}[a] = \text{tp}^{\mathfrak{A}}[a]$.
2. Let a' be any element of $A_{\{\pi\}^c}$ with $a \neq a'$. Then certainly $a' \notin A_a \subseteq A_{\pi}$, so that $\text{tp}^{\mathfrak{A}'}[a, a']$ can be different from $\text{tp}^{\mathfrak{A}}[a, a']$ only if $a \in A_{a'}$. But if $a \in A_{a'}$, then neither $\text{tp}^{\mathfrak{A}}[a, a']$ nor $\text{tp}^{\mathfrak{A}'}[a, a']$ can be a message-type. Hence, $\text{pr}_{\{\pi\}^c}^{\mathfrak{A}'}[a] = \text{pr}_{\{\pi\}^c}^{\mathfrak{A}}[a]$.
3. Suppose $a \in A \setminus B$, and let a' be any element of A with $a \neq a'$. The argument now proceeds much as for the previous condition: certainly, $a' \notin A_a = \emptyset$, so $\text{tp}^{\mathfrak{A}'}[a, a']$ can be different from $\text{tp}^{\mathfrak{A}}[a, a']$ only if $a \in A_{a'}$. But if $a \in A_{a'}$, then neither $\text{tp}^{\mathfrak{A}}[a, a']$ nor $\text{tp}^{\mathfrak{A}'}[a, a']$ can be a message-type. Hence $\text{pr}^{\mathfrak{A}'}[a] = \text{pr}^{\mathfrak{A}}[a]$.
4. Suppose $a \in B$. Equation (5) yields $\text{ct}_{\{\pi\}}^{\mathfrak{A}'}[a] = \text{ct}_{\{\pi\}}^{\mathfrak{A}}[a]$. And since B is a $\{\pi\}$ -group, $\text{ct}_{\{\pi\}}^{\mathfrak{A}}[a] = \text{ct}_{\{\pi\}}^{\mathfrak{A}}[a]$. It follows that $\text{ct}_{\{\pi\}}^{\mathfrak{A}'}[a] = \text{ct}_{\{\pi\}}^{\mathfrak{A}}[a]$.

Finally, it is immediate from Equation (5) that all elements of B have the same $\{\pi\}$ -profile in \mathfrak{A}' . \square

We now demonstrate that, if B is a Π -group in a differentiated, chromatic model of ϕ^* , then, by taking a (Π, B) -approximation to that model, we can reduce the number of Π -profiles realized by the elements of B so that it is bounded by a singly exponential function of $\|\phi^*\|$. This demonstration will occupy Lemmas 7–9. Our strategy is first to partition Π into roughly equal sets Π' and Π'' . We then recursively bound the number of Π' - and Π'' -profiles realized by elements of B , and finally align these Π' - and Π'' -profiles so as to bound the number of Π -profiles that result.

Lemma 7. *Suppose \mathfrak{A} is a chromatic model of ϕ^* . Let Π be a set of 1-types, let $B \subseteq A$ be a Π -group, and let ω be a permutation of B . Then there exists a structure \mathfrak{A}' such that \mathfrak{A}' is a (Π, B) -approximation to \mathfrak{A} , and for all $b \in B$, $\text{pr}_{\Pi}^{\mathfrak{A}'}[\omega(b)] = \text{pr}_{\Pi}^{\mathfrak{A}}[b]$.*

Proof. First, extend ω to the whole of A by setting $\omega(a) = a$ for $a \in A \setminus B$. Next, for all $b \in A$, define:

$$\omega_{b\Pi}(a) = \begin{cases} \omega(a) & \text{if } b \in A_{\Pi} \\ a & \text{otherwise;} \end{cases}$$

Thus, $\omega_{b\Pi}$ is a permutation of A (which may be the identity). Since $(\omega_{b\Pi})^{-1}$ and $(\omega^{-1})_{b\Pi}$ are the same permutation, we may unambiguously write $\omega_{b\Pi}^{-1}$. Clearly, ω fixes B setwise and $A \setminus B$ pointwise; so, therefore, does $\omega_{b\Pi}^{-1}$. Moreover, since B is a Π -group, every element of B by definition has the same 1-type, and this 1-type is either a member of Π or it is not. Hence, either $B \subseteq A_{\Pi}$ or $B \subseteq A_{\Pi^c} = A \setminus A_{\Pi}$. Thus, ω fixes both A_{Π} and A_{Π^c} setwise, and so therefore does $\omega_{b\Pi}^{-1}$.

Define the structure \mathfrak{A}' over domain A by setting, for all distinct $a, a' \in A$:

$$\text{tp}^{\mathfrak{A}'}[a, a'] = \text{tp}^{\mathfrak{A}}[\omega_{a'\Pi}^{-1}(a), \omega_{a\Pi}^{-1}(a')]. \quad (6)$$

To show that \mathfrak{A}' is well-defined, we must show first that the elements $\omega_{a'\Pi}^{-1}(a)$ and $\omega_{a\Pi}^{-1}(a')$ in each instance of (6) are distinct, and second, that the 1-type assignments implicit in the different instances of (6) do not clash. Suppose, then that $a \neq a'$; we prove that $\omega_{a'\Pi}^{-1}(a) \neq \omega_{a\Pi}^{-1}(a')$. Since the permutations $\omega_{a\Pi}^{-1}$ and $\omega_{a'\Pi}^{-1}$ fix B setwise and $A \setminus B$ pointwise, we may assume that $a, a' \in B$. We have already noted that either $B \subseteq A_{\Pi}$ or $B \subseteq A \setminus A_{\Pi}$. If $B \subseteq A_{\Pi}$, then $\omega_{a'\Pi}^{-1}(a) = \omega^{-1}(a)$ and $\omega_{a\Pi}^{-1}(a') = \omega^{-1}(a')$; if, on the other hand, $B \subseteq A \setminus A_{\Pi}$, then $\omega_{a'\Pi}^{-1}(a) = a$ and $\omega_{a\Pi}^{-1}(a') = a'$. Either way, $\omega_{a'\Pi}^{-1}(a) \neq \omega_{a\Pi}^{-1}(a')$. Next, we prove that

the 1-type assignments in (6) never clash. Since all elements of B have the same 1-type, and since ω is the identity outside B , we have, for all a, a' , $\text{tp}^{\mathfrak{A}}[\omega_{a'\Pi}^{-1}(a)] = \text{tp}^{\mathfrak{A}}[a]$; thus, $\text{tp}^{\mathfrak{A}}[\omega_{a'\Pi}^{-1}(a)]$ does not depend on a' . Hence, the 1-type assignments implicit in (6) cannot clash, and \mathfrak{A}' is indeed well-defined. In fact, this argument establishes that $\text{tp}^{\mathfrak{A}'}[a] = \text{tp}^{\mathfrak{A}}[a]$ for all $a \in A$.

We first check the numbered conditions of Definition 7 (iii) in turn. Let a be an arbitrary element of A .

1. We have just established that $\text{tp}^{\mathfrak{A}'}[a] = \text{tp}^{\mathfrak{A}}[a]$.
2. For all $b \in A_{\Pi^c}$, $\omega_{b\Pi}^{-1}(a) = a$; in particular, if $a \in A_{\Pi^c}$, then $\omega_{a\Pi}^{-1}(a) = a$. Therefore, $\omega_{a\Pi}^{-1}$ is always a permutation of $A_{\Pi^c} \setminus \{a\}$, and moreover, for all $b \in A_{\Pi^c} \setminus \{a\}$, $\text{tp}^{\mathfrak{A}'}[a, b] = \text{tp}^{\mathfrak{A}}[a, \omega_{a\Pi}^{-1}(b)]$. Thus, the list of 2-types $\text{tp}^{\mathfrak{A}'}[a, b]$ obtained as b ranges over $A_{\Pi^c} \setminus \{a\}$ is (in some order) the list of 2-types $\text{tp}^{\mathfrak{A}}[a, b']$ obtained as b' ranges over $A_{\Pi^c} \setminus \{a\}$. It follows that $\text{pr}_{\Pi^c}^{\mathfrak{A}'}[a] = \text{pr}_{\Pi^c}^{\mathfrak{A}}[a]$.
3. Suppose $a \in A \setminus B$. Then, for all $b \in A$, $\omega_{b\Pi}^{-1}(a) = a$; in particular, $\omega_{a\Pi}^{-1}(a) = a$. Therefore, $\omega_{a\Pi}^{-1}$ is a permutation of $A \setminus \{a\}$, and moreover, for all $b \in A \setminus \{a\}$, $\text{tp}^{\mathfrak{A}'}[a, b] = \text{tp}^{\mathfrak{A}}[a, \omega_{a\Pi}^{-1}(b)]$. Thus, the list of 2-types $\text{tp}^{\mathfrak{A}'}[a, b]$ obtained as b ranges over $A \setminus \{a\}$ is (in some order) the list of 2-types $\text{tp}^{\mathfrak{A}}[a, b']$ obtained as b' ranges over $A \setminus \{a\}$. It follows that $\text{pr}_{\Pi}^{\mathfrak{A}'}[a] = \text{pr}_{\Pi}^{\mathfrak{A}}[a]$.
4. For all $b \in A_{\Pi}$, $\omega_{b\Pi}^{-1}(a) = \omega^{-1}(a)$; in particular, if $a \in A_{\Pi}$, then $\omega_{a\Pi}^{-1}(a) = \omega^{-1}(a)$. Therefore, $\omega_{a\Pi}^{-1}$ is a bijection from the set $A_{\Pi} \setminus \{a\}$ to the set $A_{\Pi} \setminus \{\omega^{-1}(a)\}$, and moreover, for all $b \in A_{\Pi} \setminus \{a\}$, $\text{tp}^{\mathfrak{A}'}[a, b] = \text{tp}^{\mathfrak{A}}[\omega^{-1}(a), \omega_{a\Pi}^{-1}(b)]$. Thus, the list of 2-types $\text{tp}^{\mathfrak{A}'}[a, b]$ obtained as b ranges over $A_{\Pi} \setminus \{a\}$ is (in some order) the list of 2-types $\text{tp}^{\mathfrak{A}}[\omega^{-1}(a), b']$ obtained as b' ranges over $A_{\Pi} \setminus \{\omega^{-1}(a)\}$. It follows that

$$\text{pr}_{\Pi}^{\mathfrak{A}'}[a] = \text{pr}_{\Pi}^{\mathfrak{A}}[\omega^{-1}(a)]. \quad (7)$$

Certainly, then, we have $\text{ct}_{\Pi}^{\mathfrak{A}'}[a] = \text{ct}_{\Pi}^{\mathfrak{A}}[\omega^{-1}(a)]$. But since B is a Π -group in \mathfrak{A} , $a \in B$ implies $\text{ct}_{\Pi}^{\mathfrak{A}}[\omega^{-1}(a)] = \text{ct}_{\Pi}^{\mathfrak{A}}[a]$, whence $\text{ct}_{\Pi}^{\mathfrak{A}'}[a] = \text{ct}_{\Pi}^{\mathfrak{A}}[a]$.

We have thus established that, for all $a \in A$, $\text{pr}_{\Pi}^{\mathfrak{A}'}[a] = \text{pr}_{\Pi}^{\mathfrak{A}}[\omega^{-1}(a)]$ and $\text{pr}_{\Pi^c}^{\mathfrak{A}'}[a] = \text{pr}_{\Pi^c}^{\mathfrak{A}}[a]$. Since \mathfrak{A} is chromatic, it follows easily that \mathfrak{A}' is chromatic. Moreover, all 2-types realized in \mathfrak{A}' are realized in \mathfrak{A} . Hence, \mathfrak{A}' is a (Π, B) -approximation to \mathfrak{A} . Finally, it follows from Equation (7) that, for all $b \in B$, $\text{pr}_{\Pi}^{\mathfrak{A}'}[\omega(b)] = \text{pr}_{\Pi}^{\mathfrak{A}}[b]$. \square

Suppose \mathfrak{A} , Π and B are as in Lemma 7. That lemma then assures us that, as long as we are content to work with (Π, B) -approximations, we can permute the Π -profiles of the elements in B at will! The following lemma exploits this facility.

Lemma 8. *Suppose \mathfrak{A} is a chromatic model of ϕ^* . Let Π' , Π'' be disjoint, non-empty sets of 1-types, and let $\Pi = \Pi' \cup \Pi''$. Suppose the non-empty set $B \subseteq A$ is both a Π' -group and a Π'' -group, and hence also a Π -group. Let the number of different Π' -profiles realized in \mathfrak{A} by the elements of B be J' ; and let the number of different Π'' -profiles realized in \mathfrak{A} by the elements of B be J'' . Then there exists a (Π, B) -approximation \mathfrak{A}'' to \mathfrak{A} in which at most $J' + J'' - 1$ different Π -profiles are realized by the elements of B .*

Proof. For perspicuity, we assume first that B is finite. Enumerate B as b_1, \dots, b_I . Let the various Π' -profiles realized by at least one element of B be $\bar{v}'_1, \dots, \bar{v}'_{J'}$; and let the various Π'' -profiles realized by at least one element of B be $\bar{v}''_1, \dots, \bar{v}''_{J''}$. Since B is a Π' -group, Lemma 7 guarantees that we can obtain a (Π', B) -approximation to \mathfrak{A} in which the Π' -profiles of B are permuted at will. So let \mathfrak{A}' be a (Π', B) -approximation to \mathfrak{A} in which the Π' -profiles of the b_1, \dots, b_I fall into consecutive blocks in the sense depicted in the middle column in Fig. 1. More precisely, we have integers $0 = I_1 < I_2 < \dots < I_{J'+1} = I$ such that, for all j ($1 \leq j \leq J'$), $\text{pr}_{\Pi'}^{\mathfrak{A}'}[b_i] = \bar{v}'_j$ for i in the range $[I_j + 1, I_{j+1}]$. Since B is also a Π'' -group, we can obtain a structure \mathfrak{A}'' such that \mathfrak{A}'' is a (Π'', B) -approximation to \mathfrak{A}' in which the elements of B have Π'' -profiles likewise arranged in consecutive blocks. Since \mathfrak{A}'' is a (Π'', B) -approximation to \mathfrak{A}' and the sets Π' and Π'' are disjoint, the Π' -profiles of the elements of B will be unaffected by the transformation from \mathfrak{A}' to \mathfrak{A}'' : a typical alignment of Π' -profiles and Π'' -profiles in \mathfrak{A}'' is shown in Fig. 1. By inspection, at most $J' + J'' - 1$ Π -profiles are realized in \mathfrak{A}'' by the elements of B . From Remark 4, \mathfrak{A}'' is a (Π, B) -approximation to \mathfrak{A} , because \mathfrak{A}' is a (Π', B) -approximation to \mathfrak{A} and \mathfrak{A}'' is a (Π'', B) -approximation to \mathfrak{A}' .

The same argument applies in the case where B is infinite, with only minor modifications. Of course, the number of different Π' - and Π'' -profiles realized by the elements of B can still only be finite, but some of the resulting blocks of B may contain infinitely many entries. The matching up of these blocks so that at most $J' + J'' - 1$ Π -profiles result is routine. \square

Lemma 9. *Suppose \mathfrak{A} is a differentiated, chromatic model of ϕ^* . Let $l \geq 0$, let Π be a non-empty set of 1-types such that $|\Pi| \leq 2^l$, and let $B \subseteq A$ be a Π -group. Then there is a structure \mathfrak{A}' such that \mathfrak{A}'*

Element of B	Π' -profile in \mathfrak{A}' (and also in \mathfrak{A}'')	Π'' -profile in \mathfrak{A}''
b_1	\bar{v}'_1	\bar{v}''_1
		\bar{v}''_2
	\vdots	\vdots
		$\bar{v}''_{J''}$
b_I	$\bar{v}'_{J'}$	

Figure 1. Arrangement of Π' -profiles and Π'' -profiles in B .

is a (Π, B) -approximation to \mathfrak{A} and the number of different Π -profiles realized in \mathfrak{A}' by the elements of B is at most $2^l(M^* + 1)(C + 1)^{lm}$.

Proof. By induction on l . To aid readability, let K_l stand for $2^l(M^* + 1)(C + 1)^{lm}$. If $l = 0$, let $\Pi = \{\pi\}$. Decompose B into maximal π -patches B_1, \dots, B_H . Since \mathfrak{A} is chromatic, Remark 2 guarantees that the first M^* coordinates of any vector $\text{pr}_{\{\pi\}}^{\mathfrak{A}}[a]$ are either all zero, or else are all zero except for a single occurrence of unity. Therefore, $H \leq M^* + 1$. Now let $\mathfrak{A}_0 = \mathfrak{A}$, and for all h ($1 \leq h \leq H$), apply Lemma 6 to obtain a structure \mathfrak{A}_h such that \mathfrak{A}_h is a $(\{\pi\}, B_h)$ -approximation to \mathfrak{A}_{h-1} in which the elements of B_h all have the same $\{\pi\}$ -profile. By Remark 4, \mathfrak{A}_H is a $(\{\pi\}, B)$ -approximation to \mathfrak{A} . And because the B_h are pairwise disjoint, $1 \leq h < h' \leq H$ implies $\text{pr}^{\mathfrak{A}_{h'}}[a] = \text{pr}^{\mathfrak{A}_h}[a]$ for all $a \in B_h$. Hence, the total number of $\{\pi\}$ -profiles realized by elements of B in \mathfrak{A}_H is at most $H \leq M^* + 1 = K_0$. Thus, setting $\mathfrak{A}' = \mathfrak{A}_H$ establishes the case $l = 0$.

Now suppose $l > 0$. We may assume Π is not a singleton, since otherwise, we can employ the argument of the case $l = 0$; so let Π be partitioned into non-empty sets Π' and Π'' each of cardinality at most 2^{l-1} . Also, partition B into maximal Π' -groups B_1, \dots, B_H (say). Since B is a Π -group, the B_1, \dots, B_H will also be Π'' -groups. Moreover, since $\mathfrak{A} \models \phi^*$, the Π -count of any element in \mathfrak{A} is one of at most $(C + 1)^m$ different vectors; and since B is a Π -group, all elements of B must have the same 1-type, whence $H \leq (C + 1)^m$. Again, let $\mathfrak{A}_0 = \mathfrak{A}$, and consider the set B_1 . By inductive hypothesis, let \mathfrak{A}'_1 be a (Π', B_1) -approximation to \mathfrak{A}_0 in which at most K_{l-1} Π' -profiles are realized

by the elements of B_1 . Again, by inductive hypothesis, let \mathfrak{A}_1'' be a (Π'', B_1) -approximation to \mathfrak{A}_1' in which at most K_{l-1} Π'' -profiles are realized by the elements of B_1 . Thus, in the structure \mathfrak{A}_1'' , B_1 is a Π' -group realizing at most K_{l-1} different Π' -profiles, and also a Π'' -group realizing at most K_{l-1} different Π'' -profiles. By Lemma 8, let \mathfrak{A}_1 be a (Π, B_1) -approximation to \mathfrak{A}_1'' in which the elements of B_1 realize at most $2K_{l-1} - 1 < 2K_{l-1}$ different Π -profiles. By Remark 4, \mathfrak{A}_1 is a (Π, B_1) -approximation to \mathfrak{A}_0 . Treating the sets B_2, \dots, B_H in the same way, we obtain structures \mathfrak{A}_h ($1 \leq h \leq H$) such that, for each h in this range, \mathfrak{A}_h is a (Π, B_h) -approximation to \mathfrak{A}_{h-1} in which at most $2K_{l-1}$ different Π -profiles are realized by the elements of B_h . By Remark 4, \mathfrak{A}_H is a (Π, B) -approximation to \mathfrak{A} . And because the B_h are pairwise disjoint, $1 \leq h < h' \leq H$ implies that $\text{pr}^{\mathfrak{A}_{h'}}[a] = \text{pr}^{\mathfrak{A}_h}[a]$ for all $a \in B_h$. Hence, the total number of Π -profiles realized by elements of B in \mathfrak{A}_H is at most $2HK_{l-1} \leq 2(C+1)^m K_{l-1} = K_l$. Thus, setting $\mathfrak{A}' = \mathfrak{A}_H$ completes the induction. \square

3.3. SPARSE STRUCTURES

We are now ready to achieve the goal of this section, namely, to show that, if ϕ^* is satisfiable over some domain, then it is satisfied in a structure over that domain in which only a limited number of different sorts of element occur. Recall that μ_1, \dots, μ_{M^*} are the invertible message-types and that $\mu_{M^*+1}, \dots, \mu_M$ are the non-invertible message-types. The next definition relies on conventions established in Notation 1.

Definition 9. A *star-type* (over Σ) is a pair $\sigma = \langle \pi, \bar{v} \rangle$, where π is a 1-type over Σ and $\bar{v} = (v_1, \dots, v_M)$ is a vector over \mathbb{N} satisfying the condition that, for all j ($1 \leq j \leq M$), $v_j > 0$ implies $\text{tp}_1(\mu_j) = \pi$. We say that σ is *chromatic* if, for every 1-type π' , the sum of all the v_j ($1 \leq j \leq M^*$) such that $\text{tp}_2(\mu_j) = \pi'$ equals either 0 or 1, and equals 0 in the case $\pi' = \pi$. If $\mathfrak{A} \models \phi^*$ and $a \in A$, then $\langle \text{tp}^{\mathfrak{A}}[a], \text{pr}^{\mathfrak{A}}[a] \rangle$ is evidently a star-type, which we call the *star-type of a in \mathfrak{A}* , denoted $\text{st}^{\mathfrak{A}}[a]$. We say that the star-type σ is *realized* in \mathfrak{A} if $\sigma = \text{st}^{\mathfrak{A}}[a]$ for some $a \in A$.

It is best to think of $\text{st}^{\mathfrak{A}}[a]$ as a description of a together with its local environment. Note that, if $\langle \pi, \bar{v} \rangle$ is a star-type, and \bar{v} is not the zero-vector, then π is actually determined by \bar{v} .

Notation 3. If $\sigma = \langle \pi, \bar{v} \rangle$ is a star-type with $\bar{v} = (v_1, \dots, v_M)$, we denote π by $\text{tp}(\sigma)$ and v_j by $\sigma[j]$ for all j ($1 \leq j \leq M$).

Remark 5. If $\mathfrak{A} \models \phi^*$ and $a \in A$, then $\text{tp}(\text{st}^{\mathfrak{A}}[a]) = \text{tp}^{\mathfrak{A}}[a]$. Moreover, \mathfrak{A} is chromatic if and only if every star-type realized in \mathfrak{A} is chromatic.

Remark 6. Let σ be a chromatic star-type and let j and j' be integers between 1 and M^* (so that μ_j and $\mu_{j'}$ are invertible message-types). If $\mu_j^{-1} = \mu_{j'}$, then either $\sigma[j] = 0$ or $\sigma[j'] = 0$. In particular, if $\mu_j^{-1} = \mu_j$, then $\sigma[j] = 0$.

Definition 10. Suppose \mathfrak{A} is a model of ϕ^* , and let X be a positive integer. We say that \mathfrak{A} is X -sparse if \mathfrak{A} realizes no more than X different star-types—that is, if $|\{\text{st}^{\mathfrak{A}}[a] : a \in A\}| \leq X$.

For the next lemma, recall that s is the number of symbols in Σ , m is the number of counting predicates, and $C = \max_{1 \leq h \leq m} C_h$ (see Table I).

Lemma 10. Let $X = 4^s(16^s + 1)(C + 1)^{sm}$. If ϕ^* has a model, then it has a chromatic, differentiated, X -sparse model over the same domain.

Proof. Suppose ϕ^* has a model with domain A . By Lemmas 2 and 3, ϕ^* has a chromatic, differentiated model \mathfrak{A} with the same domain. Let Π be the set of all 1-types; thus, $|\Pi| = 2^s$. Let A_1, \dots, A_H be a list of the non-empty sets A_π , where $\pi \in \Pi$; thus, $H \leq 2^s$. The sets A_h together partition A ; moreover, since $\mathfrak{A} \models \phi^*$, each A_h ($1 \leq h \leq H$) is a Π -group. Letting $\mathfrak{A}_0 = \mathfrak{A}$, by Lemma 9, we can obtain $\mathfrak{A}_1, \dots, \mathfrak{A}_H$ such that, for all h ($1 \leq h \leq H$), \mathfrak{A}_h is a (Π, A_h) -approximation to \mathfrak{A}_{h-1} in which at most $2^s(M^* + 1)(C + 1)^{sm}$ different profiles are realized by the elements of A_h . Using by now familiar reasoning, \mathfrak{A}_H is therefore a (Π, A) -approximation to \mathfrak{A} realizing at most $H \cdot [2^s(M^* + 1)(C + 1)^{sm}] \leq 4^s(M^* + 1)(C + 1)^{sm}$ different star-types. By Lemma 5, $\mathfrak{A}_H \models \phi^*$. By Definition 7, \mathfrak{A}_H is chromatic; and since $\text{tp}^{\mathfrak{A}_H}[a] = \text{tp}^{\mathfrak{A}}[a]$ for all $a \in A$, \mathfrak{A}_H is also differentiated. Finally, $M^* \leq 2^{4s} = 16^s$. Thus, \mathfrak{A}_H is the required structure. \square

4. Deciding finite satisfiability

We continue to use the symbols in Table I with their advertised meanings. By means of Lemma 10, we have reduced the problem of determining whether ϕ^* has a (finite) model to the problem of determining whether ϕ^* has a (finite) chromatic, differentiated, X -sparse model over Σ , where $X = 4^s(16^s + 1)(C + 1)^{sm}$. Crucially, X is bounded

by a *singly* exponential function of $\|\phi^*\|$. The task of this section is to show how finite models of ϕ^* which are chromatic, differentiated and X -sparse can be encoded as data-structures satisfying certain arithmetical constraints, whose size is also bounded by a singly exponential function of $\|\phi^*\|$. The result that $\text{Fin-Sat-}\mathcal{C}^2$ is in NEXPTIME follows from this encoding.

Recall that π_1, \dots, π_L are the 1-types over Σ .

Notation 4. We write \mathcal{I} to denote the set of unordered pairs of (not necessarily distinct) integers between 1 and L . Formally: $\mathcal{I} = \{\{i, i'\} \mid 1 \leq i \leq i' \leq L\}$.

The next definition again uses Notation 1. Recall that Ξ is the set of silent 2-types over Σ .

Definition 11. A *frame* is a tuple $\mathcal{F} = (\bar{\sigma}, I, \theta)$, where $\bar{\sigma} = (\sigma_1, \dots, \sigma_N)$ is a list of pairwise distinct star-types, I is a subset of \mathcal{I} , and θ is a function $\theta : I \rightarrow \Xi$ such that, for all $\{i, i'\} \in I$ with $i \leq i'$, $\text{tp}_1(\theta(\{i, i'\})) = \pi_i$ and $\text{tp}_2(\theta(\{i, i'\})) = \pi_{i'}$. The *dimension* of \mathcal{F} is N . For Y a positive integer, \mathcal{F} is *Y -bounded* if, for all k ($1 \leq k \leq N$) and all j ($1 \leq j \leq M$), $\sigma_k[j] \leq Y$. Finally, \mathcal{F} is *chromatic* if every σ_k is chromatic.

Think of a frame $\mathcal{F} = (\bar{\sigma}, I, \theta)$ as a (putative) schematic description of a structure, where $\bar{\sigma}$ tells us which star-types are realized, I tells us which pairs of 1-types are not noisy (Definition 4), and θ selects, for each non-noisy pair of 1-types, a silent 2-type joining them. More precisely:

Definition 12. Suppose $\mathfrak{A} \models \phi^*$, and let $\mathcal{F} = (\bar{\sigma}, I, \theta)$ be a frame. We say that \mathcal{F} *describes* \mathfrak{A} if the following conditions hold:

1. $\bar{\sigma}$ is a list of all and only those star-types realized in \mathfrak{A} ;
2. I is the set of all and only those $\{i, i'\} \in \mathcal{I}$ such that π_i and $\pi_{i'}$ do not form a noisy pair in \mathfrak{A} ;
3. for each $\{i, i'\} \in I$, there exist distinct $a, a' \in A$ such that $\text{tp}^{\mathfrak{A}}[a, a'] = \theta(\{i, i'\})$.

Any model \mathfrak{A} of ϕ^* is evidently described by some (not necessarily unique) frame; and certain interesting properties of \mathfrak{A} correspond to obvious properties of the frames which describe it, as we see from the following two lemmas.

Lemma 11. *Suppose $\mathfrak{A} \models \phi^*$, and let \mathcal{F} be a frame which describes \mathfrak{A} . Then: (i) \mathfrak{A} is chromatic if and only if \mathcal{F} is chromatic; (ii) \mathfrak{A} is X -sparse if and only if \mathcal{F} has dimension at most X ; and (iii) \mathcal{F} is C -bounded.*

Proof. Immediate. □

For the next definition, recall that a 1-type π is simply a finite collection of formulas, so that $\bigwedge \pi$ denotes the conjunction of those formulas; similarly for 2-types.

Definition 13. Let $\mathcal{F} = (\bar{\sigma}, I, \theta)$ be a frame, where $\bar{\sigma} = (\sigma_1, \dots, \sigma_N)$. We write $\mathcal{F} \models \phi^*$ if the following conditions are satisfied:

1. for all k ($1 \leq k \leq N$), $\models \bigwedge \text{tp}(\sigma_k) \rightarrow \alpha$;
2. for all k ($1 \leq k \leq N$) and all j ($1 \leq j \leq M$), if $\sigma_k[j] > 0$ then $\models \bigwedge \mu_j \rightarrow \beta(x, y) \wedge \beta(y, x)$;
3. for all $\{i, i'\} \in I$, $\models \bigwedge \theta(\{i, i'\}) \rightarrow \beta(x, y) \wedge \beta(y, x)$;
4. for all k ($1 \leq k \leq N$) and all h ($1 \leq h \leq m$), the sum of all the $\sigma_k[j]$ ($1 \leq j \leq M$) such that $f_h(x, y) \in \mu_j$ equals C_h .

Remark 7. *Let $\mathcal{F} = (\bar{\sigma}, I, \theta)$ be a frame, where $\bar{\sigma} = (\sigma_1, \dots, \sigma_N)$. If $\mathcal{F} \models \phi^*$, then, for all k ($1 \leq k \leq N$), $\sum_{1 \leq j \leq M} \sigma_k[j] \leq mC$.*

Lemma 12. *Suppose $\mathfrak{A} \models \phi^*$, and let \mathcal{F} be a frame describing \mathfrak{A} . Then $\mathcal{F} \models \phi^*$.*

Proof. Almost immediate. □

However, while every model of ϕ^* is described by some frame, not every frame describes a model of ϕ^* ; and it is important for us to define a class of frames which do. Recall that π_1, \dots, π_L are the 1-types, μ_1, \dots, μ_{M^*} , the invertible message-types, and $\mu_{M^*+1}, \dots, \mu_M$, the non-invertible message-types.

Notation 5. *Let $\mathcal{F} = (\bar{\sigma}, I, \theta)$ be a frame, where $\bar{\sigma} = (\sigma_1, \dots, \sigma_N)$. If \mathcal{F} is clear from context, for integers i, k in the ranges $1 \leq i \leq L$,*

$1 \leq k \leq N$ write:

$$o_{ik} = \begin{cases} 1 & \text{if } \text{tp}(\sigma_k) = \pi_i \\ 0 & \text{otherwise;} \end{cases}$$

$$p_{ik} = \begin{cases} 1 & \text{if, for all } j \ (1 \leq j \leq M), \text{tp}_2(\mu_j) = \pi_i \text{ implies } \sigma_k[j] = 0 \\ 0 & \text{otherwise;} \end{cases}$$

$$r_{ik} = \sum_{j \in J} \sigma_k[j], \text{ where } J = \{j \mid M^* + 1 \leq j \leq M \text{ and } \text{tp}_2(\mu_j) = \pi_i\};$$

$$s_{ik} = \sum_{j \in J} \sigma_k[j], \text{ where } J = \{j \mid 1 \leq j \leq M \text{ and } \text{tp}_2(\mu_j) = \pi_i\}.$$

In addition, for integers i, j in the ranges $1 \leq i \leq L$, $1 \leq j \leq M^*$, write:

$$q_{jk} = \sigma_k[j].$$

Remark 8. Suppose $\mathfrak{A} \models \phi^*$, and let \mathcal{F} be a frame describing \mathfrak{A} . Then the symbols o_{ik} , p_{ik} , q_{jk} , r_{ik} and s_{ik} in Notation 5 have the following interpretations with respect to \mathfrak{A} :

1. $o_{ik} = 1$ just in case every element with star-type σ_k has 1-type π_i ;
2. $p_{ik} = 1$ just in case no element with star-type σ_k sends a message to any element having 1-type π_i ;
3. q_{jk} counts how many messages of (invertible) type μ_j any element having star-type σ_k sends;
4. r_{ik} is the total number of elements having 1-type π_i to which any element having star-type σ_k sends a non-invertible message; and
5. s_{ik} is the total number of elements having 1-type π_i to which any element having star-type σ_k sends a message.

With this notation in hand we can characterize a class of frames whose members are guaranteed to describe models of ϕ^* .

Definition 14. Let $\mathcal{F} = (\bar{\sigma}, I, \theta)$ be a frame, where $\bar{\sigma} = (\sigma_1, \dots, \sigma_N)$. Let $\bar{w} = (w_1, \dots, w_N)$ be a vector of positive integers. Using Notation 5, for all i ($1 \leq i \leq L$), all i' ($1 \leq i' \leq L$) and all j ($1 \leq j \leq M^*$), let:

$$u_i = \sum_{1 \leq k \leq N} o_{ik} w_k \quad v_j = \sum_{1 \leq k \leq N} q_{jk} w_k \quad x_{ii'} = \sum_{1 \leq k \leq N} o_{ik} p_{i'k} w_k.$$

We say that \bar{w} is a *solution* of \mathcal{F} if the following conditions are satisfied for all i ($1 \leq i \leq L$), all i' ($1 \leq i' \leq L$), all j ($1 \leq j \leq M^*$) and all k ($1 \leq k \leq N$):

(C1) $v_j = v_{j'}$, where j' is such that $\mu_j^{-1} = \mu_{j'}$;

(C2) $s_{ik} \leq u_i$;

(C3) $u_i \leq 1$ or $u_i > Z$;

(C4) if $o_{ik} = 1$, then either $u_i > 1$ or $r_{i'k} \leq x_{i'i}$;

(C5) if $\{i, i'\} \notin I$, then either $u_i \leq 1$ or $u_{i'} \leq 1$;

(C6) if $\{i, i'\} \notin I$ and $o_{ik} = 1$, then $r_{i'k} \geq x_{i'i}$.

Remark 9. Suppose \mathfrak{A} is a finite model of ϕ^* , and let $\mathcal{F} = (\bar{\sigma}, I, \theta)$ be a frame describing \mathfrak{A} . For all k ($1 \leq k \leq N$), let w_k be the number of elements of A having star-type σ_k in \mathfrak{A} . In that case, the symbols u_i , v_j and $x_{i'i'}$ in Definition 14 have the following interpretations with respect to \mathfrak{A} :

1. u_i is the number of elements $a \in A$ such that $\text{tp}^{\mathfrak{A}}[a] = \pi_i$;
2. v_j is the number of pairs $\langle a, b \rangle \in A^2$ such that $a \neq b$ and $\text{tp}^{\mathfrak{A}}[a, b] = \mu_j$;
3. $x_{i'i'}$ is the number of elements $a \in A$ such that $\text{tp}^{\mathfrak{A}}[a] = \pi_i$ and a does not send a message to any element having 1-type $\pi_{i'}$.

The following Lemma shows that Definition 14 is not too stringent for the structures that interest us.

Lemma 13. Suppose \mathfrak{A} is a finite, differentiated model of ϕ^* . Let $\mathcal{F} = (\bar{\sigma}, I, \theta)$ be a frame describing \mathfrak{A} . Then \mathcal{F} has a solution.

Proof. Let $\bar{\sigma} = (\sigma_1, \dots, \sigma_N)$, and let $w_k = |\{a \in A : \text{st}^{\mathfrak{A}}[a] = \sigma_k\}|$ for all k ($1 \leq k \leq N$). We show that $\bar{w} = (w_1, \dots, w_N)$ is a solution of \mathcal{F} . In doing so, we make free use of Remarks 8 and 9. Note that, by construction, the w_1, \dots, w_N are all positive.

C1: If $\mu_j^{-1} = \mu_{j'}$, then the sets $\{\langle a, b \rangle \mid a \neq b \text{ and } \text{tp}^{\mathfrak{A}}[a, b] = \mu_j\}$ and $\{\langle a, b \rangle \mid a \neq b \text{ and } \text{tp}^{\mathfrak{A}}[a, b] = \mu_{j'}\}$ can obviously be put in 1-1 correspondence, namely: $\langle a, b \rangle \mapsto \langle b, a \rangle$. But the cardinalities of these sets are v_j and $v_{j'}$, respectively.

C2: Since \mathcal{F} describes \mathfrak{A} , any element of A having star-type σ_k sends a message to exactly s_{ik} elements having 1-type π_i . But u_i is the number of elements of A having 1-type π_i . Since σ_k is realized in \mathfrak{A} , $s_{ik} \leq u_i$.

C3: Immediate given that \mathfrak{A} is differentiated.

C4: If $o_{ik} = 1$ and $u_i \leq 1$, then $u_i = 1$, so that \mathfrak{A} contains exactly one element with 1-type π_i ; moreover, this element has star-type σ_k . Denote this element by a . Thus, a sends a non-invertible message to exactly $r_{i'k}$ elements with 1-type $\pi_{i'}$. Clearly, none of these elements sends a message back to a (since otherwise a 's message to it would be invertible), so that there exist at least $r_{i'k}$ elements with 1-type $\pi_{i'}$ which do not send a message to a . But since a is the only element with 1-type π_i , there exist at least $r_{i'k}$ elements with 1-type $\pi_{i'}$ which do not send a message to any element having 1-type π_i . In other words, $r_{i'k} \leq x_{i'i}$.

C5: Since \mathcal{F} describes \mathfrak{A} , $\{i, i'\} \notin I$ implies that π_i and $\pi_{i'}$ form a noisy pair in \mathfrak{A} . In that case, by Lemma 4, either $u_i \leq 1$ or $u_{i'} \leq 1$.

C6: Since \mathcal{F} describes \mathfrak{A} , $\{i, i'\} \notin I$ implies that π_i and $\pi_{i'}$ form a noisy pair in \mathfrak{A} . Now if $o_{ik} = 1$, there exists at least one element a having 1-type π_i and star-type σ_k . Moreover, there are $x_{i'i}$ elements having 1-type $\pi_{i'}$ which do not send a message to any element having 1-type π_i , and hence at least $x_{i'i}$ elements having 1-type $\pi_{i'}$ which do not send a message to a . Therefore, a sends a message (in fact, a non-invertible message) to all of these elements. But since a has star-type σ_k , a sends a non-invertible message to exactly $r_{i'k}$ elements having 1-type $\pi_{i'}$. Thus, $r_{i'k} \geq x_{i'i}$. \square

We now prove a converse of Lemma 13.

Lemma 14. *Let \mathcal{F} be a chromatic frame. If \mathcal{F} has a solution and $\mathcal{F} \models \phi^*$, then there exists a finite structure \mathfrak{A} such that $\mathfrak{A} \models \phi^*$.*

Proof. Let $\mathcal{F} = (\bar{\sigma}, I, \theta)$, let $\bar{\sigma} = (\sigma_1, \dots, \sigma_N)$, and let $\bar{w} = (w_1, \dots, w_N)$ be a solution of \mathcal{F} . In the sequel, we use the symbols $o_{ik}, p_{ik}, q_{jk}, r_{ik}$ and s_{ik} (with indices in the appropriate ranges), as specified in Notation 5, and the symbols u_i, v_j and $x_{i'i'}$ (again, with indices in the appropriate ranges), as specified in Definition 14. Hence, the conditions **C1–C6** of Definition 14 hold.

For every k ($1 \leq k \leq N$), let A_k be a set of cardinality w_k , and let A be the disjoint union of the A_k . Think of A_k as the set of elements which ‘want’ to have star-type σ_k . In addition, we define for all i ($1 \leq i \leq L$), all i' ($1 \leq i' \leq L$) and all j ($1 \leq j \leq M^*$):

$$\begin{aligned} U_i &= \bigcup \{A_k \mid 1 \leq k \leq N \text{ and } o_{ik} = 1\} \\ X_{i'i'} &= \bigcup \{A_k \mid 1 \leq k \leq N \text{ and } o_{ik} p_{i'k} = 1\} \\ V_j &= \bigcup \{A_k \mid 1 \leq k \leq N \text{ and } q_{jk} = 1\}. \end{aligned}$$

Since \mathcal{F} is chromatic, $q_{jk} \leq 1$ for all j ($1 \leq j \leq M^*$) and all k ($1 \leq k \leq N$). Thus, for all i, i' and j in the appropriate ranges:

$$|U_i| = u_i; \quad |X_{ii'}| = x_{ii'}; \quad |V_j| = v_j.$$

Think of U_i as the set of elements which ‘want’ to have 1-type π_i , $X_{ii'}$ as the set of elements in U_i which do not ‘want’ to send a message to any element in $U_{i'}$, and V_j as the set of elements which ‘want’ to send an (invertible) message of type μ_j to some other element. We remark that $A_k \subseteq U_i$ if and only if $\text{tp}(\sigma_k) = \pi_i$. Moreover, for all j ($1 \leq j \leq M^*$), if $V_j \neq \emptyset$, there exists a unique i ($1 \leq i \leq L$) such that $V_j \subseteq U_i$ —namely, that i such that $\text{tp}_1(\mu_j) = \pi_i$. We now convert the domain A into a structure \mathfrak{A} in four steps.

Step 1 (Interpreting the unary predicates and diagonals of binary predicates): For every k ($1 \leq k \leq N$) and every $a \in A_k$, set $\text{tp}^{\mathfrak{A}}[a] = \text{tp}(\sigma_k)$. At the end of this step, we have, for every i ($1 \leq i \leq L$) and every $a \in U_i$, $\text{tp}^{\mathfrak{A}}[a] = \pi_i$.

Step 2 (Fixing the invertible message-types): For every j ($1 \leq j \leq M^*$), let j' be such that $\mu_j^{-1} = \mu_{j'}$. By **C1**, V_j and $V_{j'}$ are equinumerous. If $j' > j$, pick some 1–1 correspondence between V_j and $V_{j'}$; and for every $a \in V_j$, set $\text{tp}^{\mathfrak{A}}[a, a'] = \mu_j$, where a' is the element of $V_{j'}$ corresponding to $a \in V_j$. This completes Step 2. We must show that these assignments are meaningful, do not clash with Step 1, and do not clash with each other. Suppose then that the assignment $\text{tp}^{\mathfrak{A}}[a, a'] = \mu_j$ is made, and that $\mu_j^{-1} = \mu_{j'}$. Thus, $a \in V_j$ and $a' \in V_{j'}$. To show that the assignment is meaningful, we must prove that $a \neq a'$. For contradiction, suppose $a = a'$, and let k be such that $a \in A_k$. But then $\sigma_k[j] > 0$ and $\sigma_k[j'] > 0$, which is impossible by Remark 6. To show that the assignment does not clash with Step 1, suppose $\mu_j^{-1} = \mu_{j'}$, and let i, i' be such that $V_j \subseteq U_i$ and $V_{j'} \subseteq U_{i'}$. As observed above, $\pi_i = \text{tp}_1(\mu_j)$ and $\pi_{i'} = \text{tp}_1(\mu_{j'}) = \text{tp}_2(\mu_j)$, which conforms to the assignments in Step 1. To show that these assignments do not clash with each other, it suffices to prove that, if $a \in V_j \cap V_h$, $a' \in V_{j'} \cap V_{h'}$, $\mu_j^{-1} = \mu_{j'}$ and $\mu_h^{-1} = \mu_{h'}$, then $j = h$. Suppose then that the antecedent of this conditional holds; let k and k' be such that $a \in A_k$ and $a' \in A_{k'}$. Then $\sigma_{k'}[j'] > 0$ and $\sigma_{k'}[h'] > 0$. Since $\sigma_{k'}$ is a star-type, $\text{tp}_1(\mu_{j'}) = \text{tp}_1(\mu_{h'})$, whence $\text{tp}_2(\mu_j) = \text{tp}_2(\mu_h)$. But $\sigma_k[j] > 0$ and $\sigma_k[h] > 0$, and since σ_k is a *chromatic* star-type, $j = h$. Note that, if $\mu_j^{-1} = \mu_j$, then $V_j = \emptyset$ by Remark 6. Thus, at the end of Step 2, for every element $a \in A$ and every j ($1 \leq j \leq M^*$), a sends a (unique) message of type μ_j to some other element if and only if $a \in V_j$. That is: for all k ($1 \leq k \leq N$), all $a \in A_k$, and all j ($1 \leq j \leq M^*$), there are exactly $\sigma_k[j]$ elements




Figure 2. Dealing with non-invertible messages between U_i and $U_{i'}$.

$a' \in A$ such that $a \neq a'$ and $\text{tp}^{\mathfrak{A}}[a, a'] = \mu_j$. We make one further observation before proceeding. Suppose that $\text{tp}^{\mathfrak{A}}[a, a']$ is assigned in this step and that $a \in U_i$; we claim that $a' \notin X_{i'i}$ for any i' . To see this, suppose $a \in V_j \subseteq U_i$ and $a' \in A_{k'} \subseteq V_{j'}$, with $\mu_j^{-1} = \mu_{j'}$. Then $\text{tp}_1(\mu_j) = \text{tp}_2(\mu_{j'}) = \pi_i$. But then $\sigma_{k'}[j'] > 0$, whence $p_{ik'} = 0$, whence $a' \notin X_{i'i}$. This observation will be useful in Step 3.

Step 3 (Fixing the non-invertible message-types): Let i and i' be such that $1 \leq i \leq i' \leq L$. We fix all the non-invertible messages sent, in either direction, between U_i and $U_{i'}$. By **C3**, either $u_i \leq 1$ or $u_i > Z$; similarly, either $u_{i'} \leq 1$ or $u_{i'} > Z$. We consider five cases.

Case 1: $u_i = 0$. In this case, there are no elements of U_i and hence no 2-type assignments to be made between elements of U_i and elements of $U_{i'}$. Note that, by **C2**, $s_{ik} = 0$ for all k ($1 \leq k \leq N$), whence $\sigma_k[j] = 0$ for all k ($1 \leq k \leq N$) and for all j ($1 \leq j \leq M$) such that $\text{tp}_2(\mu_j) = \pi_i$. (Intuitively, no element of A —and in particular of $U_{i'}$ —‘wants’ to send a message to an element with 1-type π_i anyway.)

Case 2: $u_i = 1$. The situation is illustrated in the left-hand diagram of Fig. 2. Let a be the sole element of U_i , and let k be such that $a \in A_k$. We deal first with the assignment of non-invertible messages sent from $U_{i'}$ to $U_i = \{a\}$. Consider any $a' \in A_{k'} \subseteq U_{i'}$. By **C2**, $s_{ik'} \leq 1$; hence there is at most one value of j in the range $1 \leq j \leq M$ such that $\text{tp}_2(\mu_j) = \pi_i$ and $\sigma_{k'}[j] > 0$. Suppose then that such a j exists. Again, since $s_{ik'} \leq 1$, $\sigma_{k'}[j] = 1$. If $j \leq M^*$, then this message has already been dealt with in Step 2; so we may assume $M^* + 1 \leq j \leq M$. It follows

from **C4** that $i \neq i'$. Hence $a \neq a'$, so that we may set $\text{tp}^{\mathfrak{A}}[a', a] = \mu_j$. Since $\text{tp}_1(\mu_j) = \pi_{i'}$ and $\text{tp}_2(\mu_j) = \pi_i$, this assignment does not clash with Step 1. Observe also that, just as in Step 2, if this assignment is made, we have, by definition, $p_{ik'} = 0$, so that $a' \notin X_{i'i}$. By carrying out the same procedure for all $a' \in U_{i'}$, we complete the assignment of non-invertible messages sent from $U_{i'}$ to U_i . It remains to deal with the non-invertible messages sent from $U_i = \{a\}$ to $U_{i'}$. Remembering that $a \in A_k$, **C4** ensures the existence of a subset $R \subseteq X_{i'i}$ such that $|R| = r_{i'k}$. For each j ($M^* + 1 \leq j \leq M$), if $\text{tp}_2(\mu_j) = \pi_{i'}$, select $\sigma_k[j]$ fresh elements a' of R , and make the assignment $\text{tp}^{\mathfrak{A}}[a, a'] = \mu_j$. (There are enough such elements by the definition of $r_{i'k}$.) These assignments clearly do not clash with those made in Step 1. Moreover, we have observed that $\text{tp}^{\mathfrak{A}}[a, a']$ has previously been assigned (either in this step or in Step 2) only if $a' \notin X_{i'i}$. Thus, these assignments do not clash with those made earlier in this step or those made in Step 2.

Case 3: $u_{i'} = 0$ and $u_i > Z$. Symmetrical to Case 1.

Case 4: $u_{i'} = 1$ and $u_i > Z$. Symmetrical to Case 2.

Case 5: $u_i > Z$ and $u_{i'} > Z$. Since $Z = (mC + 1)^2 > 3mC$, partition U_i into three sets U_{i0}, U_{i1}, U_{i2} , each containing at least mC elements; and similarly for $U_{i'}$. Suppose $a \in U_i$. Then for some h ($0 \leq h < 3$), $a \in U_{ih}$. Let k be such that $a \in A_k$, and let $h' = h + 1 \pmod{3}$. For all j ($M^* + 1 \leq j \leq M$), select $\sigma_k[j]$ fresh elements a' of $U_{i'h'}$ such that $\text{tp}^{\mathfrak{A}}[a, a']$ was not assigned in Step 2, and set $\text{tp}^{\mathfrak{A}}[a, a'] = \mu_j$. By Remark 7, $\sum_{1 \leq j \leq M} \sigma_k[j] \leq mC$; and since $|U_{i'h'}| \geq mC$, we never run out of fresh elements to select. In this way, we deal with all messages sent from U_i to $U_{i'}$; the messages sent from $U_{i'}$ to U_i are dealt with symmetrically. It is obvious that these assignments do not clash with Step 1 or Step 2; and the fact that $h' = h + 1 \pmod{3}$, ensures that they do not clash with each other (even if $i = i'$), as is evident from the right-hand diagram of Fig. 2.

Performing these assignments for all pairs i, i' such that $1 \leq i \leq i' \leq L$ completes Step 3. At the end of Step 3, then, for all k ($1 \leq k \leq N$), all $a \in A_k$, and all j ($1 \leq j \leq M$), there are exactly $\sigma_k[j]$ elements $a' \in A$ such that $a \neq a'$ and $\text{tp}^{\mathfrak{A}}[a, a'] = \mu_j$.

Step 4 (Fixing the silent 2-types): Finally, we use the components I and θ of $\mathcal{F} = (\bar{\sigma}, I, \theta)$ to deal with all the remaining 2-types in \mathfrak{A} . Let a, a' be distinct elements of A such that $\text{tp}^{\mathfrak{A}}[a, a']$ has not yet been assigned. Let i, i', k, k' be such that $a \in A_k \subseteq U_i$ and $a' \in A_{k'} \subseteq U_{i'}$, and assume, without loss of generality, that $i \leq i'$. We claim that $\{i, i'\} \in I$. For suppose otherwise. By **C5**, we have either $u_i = 1$ or $u_{i'} = 1$. Assume

the former. Now, if $p_{ik'} = 0$, then there is some j' ($1 \leq j' \leq M$) such that $\sigma_{k'}[j'] > 0$ and $\text{tp}_2(\mu_{j'}) = \pi_i$, whence—bearing in mind that a is the unique element of U_i — $\text{tp}^{\mathfrak{A}}[a, a']$ will certainly have been assigned in Step 2 (if $\mu_{j'}$ is an invertible message-type) or in Step 3 Case 2 (if $\mu_{j'}$ is a non-invertible message-type), contradicting the fact that $\text{tp}^{\mathfrak{A}}[a, a']$ is unassigned. Thus, $p_{ik'} = 1$, and hence $o_{i'k'} p_{ik'} = 1$. That is: $a' \in X_{i'i}$. But $|X_{i'i}| = x_{i'i}$. And by **C6**, $x_{i'i} \leq r_{i'k}$. Yet in Step 3 (Case 2), $r_{i'k}$ elements of $X_{i'i}$ were chosen to receive messages from a . Hence a' must be among these elements, again contradicting the fact that $\text{tp}^{\mathfrak{A}}[a, a']$ is unassigned. The case where $u_{i'} \leq 1$ proceeds symmetrically. Thus, we have established that, if $\text{tp}^{\mathfrak{A}}[a, a']$ has not yet been assigned, then $\{i, i'\} \in I$, so that we can make the assignment $\text{tp}^{\mathfrak{A}}[a, a'] = \theta(\{i, i'\})$. Since $\text{tp}_1(\theta(\{i, i'\})) = \pi_i$ and $\text{tp}_2(\theta(\{i, i'\})) = \pi_{i'}$, there is no clash with Step 1. Evidently, we can proceed in this way until all remaining 2-types have been assigned. Moreover, since each $\theta(\{i, i'\})$ is silent, this step does not spoil the work of Steps 2–3: we still have that, for all k ($1 \leq k \leq N$), all $a \in A_k$, and all j ($1 \leq j \leq M$), there are exactly $\sigma_k[j]$ elements $a' \in A$ such that $a \neq a'$ and $\text{tp}^{\mathfrak{A}}[a, a'] = \mu_j$.

This completes the construction of \mathfrak{A} . The only 1-types realized in \mathfrak{A} are the 1-types $\text{tp}(\sigma_k)$ (where $1 \leq k \leq N$). The only message-types realized in \mathfrak{A} are those μ_j such that $\sigma_k[j] > 0$ for some k . And the only silent 2-types realized in \mathfrak{A} are the $\theta(\{i, i'\})$ for $\{i, i'\} \in I$. Since $\mathcal{F} \models \phi^*$, we have $\mathfrak{A} \models \forall x \alpha \wedge \forall x \forall y (\beta \vee x \approx y)$. Moreover, for all k ($1 \leq k \leq N$), and for all $a \in A_k$, $\text{st}^{\mathfrak{A}}[a] = \sigma_k$. Since $\mathcal{F} \models \phi^*$, we have, for all $a \in A$, $\text{ct}^{\mathfrak{A}}[a] = (C_1, \dots, C_m)$; in other words, $\mathfrak{A} \models \bigwedge_{1 \leq h \leq m} \forall x \exists =_{C_h} y (f_h(x, y) \wedge x \not\approx y)$. Hence, $\mathfrak{A} \models \phi^*$. \square

Lemmas 13 and 14 in effect reduce the task of determining whether ϕ^* is finitely satisfiable to that of determining whether certain frames \mathcal{F} have a solution. We can now employ a standard result to bound the complexity of determining whether a given frame has a solution.

Lemma 15. *Let \mathcal{F} be a Y -bounded, N -dimensional frame over Σ . Then \mathcal{F} has a solution if and only if it has a solution \bar{w} such that every component of \bar{w} is bounded by some (fixed) singly exponential function of the quantity $L + M^* + N + \log Y + \log Z$.*

Proof. This follows immediately from the well-known result (Papadimitriou [5]) that, if an integer programming problem has a solution at all, then it has a solution all of whose components are bounded by a singly exponential function of the size of the problem (encoded in the obvious way). For the conditions **C1–C6** in Definition 14 amount to a disjunction of integer programming problems whose sizes are all bounded by a polynomial function of $L + M^* + N + \log Y + \log Z$. \square

Theorem 1. *The problem Fin-Sat- \mathcal{C}^2 is in NEXPTIME.*

Proof. By Lemma 1, it suffices to show that the finite satisfiability of any formula ϕ^* of the form (1) can be decided non-deterministically in time bounded by a singly exponential function of $\|\phi^*\|$. Using the symbols in Table I with the advertised interpretations, let $X = 4^s(16^s + 1)(C + 1)^{sm}$.

We claim that ϕ^* is finitely satisfiable if and only if there exists a chromatic, C -bounded frame \mathcal{F} over Σ of dimension $N \leq X$, such that \mathcal{F} has a solution and $\mathcal{F} \models \phi^*$. For suppose ϕ^* is finitely satisfiable. By Lemma 10, ϕ^* has a finite, chromatic, differentiated, X -sparse model \mathfrak{A} interpreting Σ . Let \mathcal{F} be a frame over Σ describing \mathfrak{A} . By Lemma 11, \mathcal{F} is chromatic, of dimension $N \leq X$ and C -bounded. By Lemma 12, $\mathcal{F} \models \phi^*$. Finally, by Lemma 13, \mathcal{F} has a solution. Conversely, suppose \mathcal{F} is a chromatic frame over Σ such that \mathcal{F} has a solution and $\mathcal{F} \models \phi^*$. Then Lemma 14 guarantees that ϕ^* is finitely satisfiable.

By Lemma 15, then, ϕ^* is finitely satisfiable if and only if there exists a chromatic, C -bounded frame \mathcal{F} over Σ of dimension $N \leq X$ and a vector \bar{w} of positive integers bounded by some doubly exponential function of $\|\phi^*\|$, such that $\mathcal{F} \models \phi^*$ and \bar{w} is a solution of \mathcal{F} . Using the standard binary encoding of integers, it is easy to write down \mathcal{F} and \bar{w} in a number of bits bounded by a singly exponential function of $\|\phi^*\|$, and to check whether they satisfy the requisite conditions in time bounded by a singly exponential function of $\|\phi^*\|$. \square

The above proof yields a small model property for finitely satisfiable \mathcal{C}^2 -formulas:

Corollary 1. *Let ϕ be a formula of \mathcal{C}^2 . If ϕ is finitely satisfiable, then it is satisfiable in a structure of size bounded by a doubly exponential function of $\|\phi\|$.*

Proof. The structure built in Lemma 14 from \mathcal{F} and its solution \bar{w} has domain of cardinality $w_1 + \dots + w_N$. \square

Notice that the complexity result of Theorem 1 is better than one might naïvely expect on the basis of the small model property of Corollary 1. Nevertheless, the bound of Corollary 1 is optimal in the sense that there is a sequence $\{\phi_i\}$ of finitely satisfiable formulas of \mathcal{C}^2 whose size grows as a polynomial function of i , but whose smallest satisfying structures grow as a doubly exponential function of i (Grädel and Otto [2], p. 317).

5. Deciding satisfiability

Having established that $\text{Fin-Sat-}\mathcal{C}^2$ is in NEXPTIME, we now turn our attention to $\text{Sat-}\mathcal{C}^2$. In fact, there is almost no further work to do.

Notation 6. Let \mathbb{N}^* denote the set $\mathbb{N} \cup \{\aleph_0\}$. We extend the ordering $>$ and the arithmetic operations $+$ and \cdot from \mathbb{N} to \mathbb{N}^* in the obvious way. Specifically, we define $\aleph_0 > n$ for all $n \in \mathbb{N}$; we define $\aleph_0 + \aleph_0 = \aleph_0 \cdot \aleph_0 = \aleph_0$ and $0 \cdot \aleph_0 = \aleph_0 \cdot 0 = 0$; we define $n + \aleph_0 = \aleph_0 + n = \aleph_0$ for all $n \in \mathbb{N}$; and we define $n \cdot \aleph_0 = \aleph_0 \cdot n = \aleph_0$ for all $n \in \mathbb{N}$ such that $n > 0$. Under this extension, $>$ remains a total order, and $+$, \cdot remain associative and commutative.

A system of linear equalities and inequalities defining an integer programming problem can of course be re-interpreted so that solutions are sought not over \mathbb{N} but over \mathbb{N}^* . (We always assume that the coefficients occurring in such problems are in \mathbb{N} .) As an example, the single inequality $x_1 \geq x_1 + 1$ has no solutions over \mathbb{N} , but it does have a solution over \mathbb{N}^* , namely, $x_1 = \aleph_0$.

Lemmas 2–10 apply to both finite and infinite structures. Furthermore, the definition of a frame and its relationship to the models of ϕ^* it describes make no reference to the cardinalities of those models, and Lemmas 11–12 again apply generally. Definition 14 requires modification, however.

Definition 15. Let Σ and $\mathcal{F} = (\bar{\sigma}, I, \theta)$ be as in Definition 14. Let $\bar{w} = (w_1, \dots, w_N)$ be a vector of non-zero elements of \mathbb{N}^* . We say that \bar{w} is an *extended solution* of \mathcal{F} if \bar{w} satisfies the conditions of Definition 14, with the arithmetic interpreted over \mathbb{N}^* as specified in Notation 6.

We must check that the obvious analogues of Lemmas 13 and 14 hold:

Lemma 16. *Suppose \mathfrak{A} is a differentiated model of ϕ^* . Let $\mathcal{F} = (\bar{\sigma}, I, \theta)$ be a frame describing \mathfrak{A} . Then \mathcal{F} has an extended solution.*

Lemma 17. *Let \mathcal{F} be a chromatic frame. If \mathcal{F} has an extended solution and $\mathcal{F} \models \phi^*$, then there exists a structure \mathfrak{A} such that $\mathfrak{A} \models \phi^*$.*

The proofs are exactly the same as in the finite case. Note that the variables u_i , v_j and $x_{ii'}$ as well as the w_k may now take the value \aleph_0 ; by contrast, the coefficients o_{ik} , p_{ik} , q_{jk} , r_{ik} and s_{ik} remain finite. Remark 9 generalizes unproblematically to countably infinite structures, so that the quantities u_i , v_j and $x_{ii'}$ mentioned in Definition 14 continue to have their familiar interpretations. The proofs of Lemmas 16 and 17 then proceed exactly as for Lemmas 13 and 14.

There is one final hurdle to overcome. The proof of Lemma 15 used a well-known result bounding solutions of integer programming problems. Since we are now dealing with \mathbb{N}^* -programming problems, we need the following extension of that result.

Lemma 18. *Let Φ be a finite set of linear inequalities of the form*

$$a_0 + a_1x_1 + \cdots + a_nx_n \leq b_0 + b_1x_1 + \cdots + b_nx_n$$

in variables x_1, \dots, x_n . Here, all coefficients are assumed to be in \mathbb{N} . We take the size of Φ , denoted $\|\Phi\|$, to be measured in the usual way, assuming binary encoding of integers. If Φ has a solution over \mathbb{N}^ , then Φ has a solution over \mathbb{N}^* such that all finite values are bounded by some (fixed) singly exponential function of $\|\Phi\|$.*

Proof. Suppose that Φ has a solution over \mathbb{N}^* . Re-order the variables if necessary so that this solution has the form $\bar{w}\bar{\aleph}_0$, with $\bar{w} = w_1, \dots, w_k \in \mathbb{N}^k$ for some k ($0 \leq k \leq n$) and $\bar{\aleph}_0$ an $(n - k)$ -tuple of \aleph_0 s. Say that an inequality in Φ *does not involve* the variable x_i if the corresponding coefficients a_i and b_i are both zero. Let Ψ be the set of inequalities in Φ involving none of the x_{k+1}, \dots, x_n . Thus, Ψ , considered as a problem in variables x_1, \dots, x_k , has a solution \bar{w} over \mathbb{N} , whence it has a solution \bar{w}' bounded by some singly exponential function of $\|\Psi\|$ (and hence of $\|\Phi\|$). But then it is easy to see that $\bar{w}'\bar{\aleph}_0$ is a solution of Φ . \square

Theorem 2. *The problem Sat- \mathcal{C}^2 is in NEXPTIME.*

Proof. Exactly as for Theorem 1, noting that, by Lemma 18, extended solutions for a C -bounded frame \mathcal{F} over Σ of dimension $N \leq 4^s(16^s + 1)(C + 1)^{sm}$ can be written down and checked in time bounded by an exponential function of $\|\phi^*\|$. \square

Obviously, there is no interesting small model property for satisfiable \mathcal{C}^2 -formulas along the lines of Corollary 1. However, we have the next best thing:

Corollary 2. *Let ϕ be a formula of \mathcal{C}^2 . Then there exist integers X and W , with X bounded by a singly exponential function of $\|\phi\|$ and W by a doubly exponential function of $\|\phi\|$, such that, if ϕ is satisfiable, then it has an X -sparse model in which every star-type is realized either infinitely often or at most W times.*

References

1. Grädel, E. and M. Otto: 1999, 'On logics with two variables'. *Theoretical Computer Science* **224**(1–2), 73–113.

2. Grädel, E., M. Otto, and E. Rosen: 1997, 'Two-variable logic with counting is decidable'. In: *Proceedings of the 12th IEEE Symposium on Logic in Computer Science*. pp. 306–317.
3. Pacholski, L., W. Szwoast, and L. Tendera: 1997, 'Complexity of two-variable logic with counting'. In: *Proceedings of the 12th IEEE Symposium on Logic in Computer Science*. pp. 318–327.
4. Pacholski, L., W. Szwoast, and L. Tendera: 1999, 'Complexity results for first-order two-variable logic with counting'. *SIAM Journal on Computing* **29**(4), 1083–1117.
5. Papadimitriou, C. H.: 1981, 'On the complexity of integer programming'. *Journal of the Association for Computing Machinery* **28**(4), 765–768.